

We claim:

- subj 50/*
1. A multiprocessor computer system comprising:
first, second and other processing nodes, each including at least one processor;
a communication pathway connecting said nodes and including a central hardware
device;
a shared, distributed system memory, a portion of said shared system memory being
coupled to said processors and to said communication pathway; and
wherein said central hardware device communicates requests from said first node to
~~said second node but to said other node.~~
 - 10 2. The system of claim 1 wherein further said nodes each includes local memory which
is directly addressable both locally by the node associated with said memory and remotely by
any other node.
 - 15 3. The system of claim 2 wherein further said central hardware device stores
information for determining which nodes or processors are storing copies of one or more
identified data elements in each said node's local memory.
 - 20 4. A multiprocessor computer system capable of being partitioned into one or more
independently functioning processing systems comprising:
a plurality of processing nodes;
a shared, distributed system memory;
a communications pathway which interconnects said plurality of processing nodes, each
node capable of operating independently, wherein each one of said processing nodes includes
at least one processor and a portion of said shared system memory coupled to said processor
and said communication pathway, and wherein said communication pathway is comprised of a
central hardware device including tag and address means to communicate the identification of
25

data transactions being processed through the system connected to said plurality of processing nodes, said tag and address means including means to store information related to the identification of a partition to which said data is associated;

5 means to define one or more partitions to which one or more of said nodes are a member; and

 means to associate each data transaction processed in the system to direct data only to such nodes defined as within the same partition of said data.

10 5. The multiprocessor system of claim 4 wherein further each node in the system includes memory which is accessible locally by the node associated with said memory and remotely by any other node.

 6. The multiprocessor system of claim 5 wherein further said central hardware device stores information for determining which nodes are storing copies of data which is associated with a given partition.

15 7. The multiprocessor computer system of claim 6 wherein said means to associate each data transaction processed in the system to direct data only to such nodes defined as within the same partition of said data includes dispatch means which interconnect with said ports but allow communications with said ports only for data transactions defined for a partition for which said port is a member.

20 8. A multiprocessor computer system capable of being partitioned into one or more independently functioning processing systems comprising:

 a plurality of processing nodes, each node capable of operating independently;
 a shared, distributed system memory;
 a communications pathway which interconnects said plurality of processing nodes,
 wherein each one of said processing nodes includes at least one processor and a portion of said

shared system memory coupled to said processor and said communication pathway, wherein
said communication pathway is comprised of a central hardware device including tag and
address means to communicate the identification of data transactions being processed through
the system connected to said plurality of processing nodes, said tag and address means
including a first register means to store information related to the identification of a partition to
which said data is associated, and a second register means to define at most one partition to
which one or more of said nodes are a member; and

means to associate each data transaction processed in the system such as to direct data
only to such nodes defined as within the same partition of said data.

10 9. The system of claim 8 wherein said tag and address means further includes a
registers means associated with each node, said register means separately connected through
said communications pathway only to one of said nodes, whereby each of said register means is
connected to only one of said nodes, wherein each said register means accepts and forwards
data transactions only identified as valid for the partition which said register is connected.

15 10. The system of claim 8 wherein further said tag and address means includes
dispatch means which restrict output of data transactions to only nodes which are defined within
the partition with which the data is identified, and a second target node which is the destination
of the requested data, wherein further said central hardware device transmits requests from the
said requesting node to the said target node but not to any other node.

20 11. In a multiprocessor computer system comprising a plurality of processing nodes; a
shared, distributed system memory; and a communication pathway connecting said processing
nodes; wherein each one of said processing nodes includes at least one processor; and a
portion of said shared system memory coupled to said processor and said communication
pathway; said communications pathway comprised of communications ports each dedicated to

communicating with one of said processing nodes; a method for partitioning the resources of the system into two or more partitions; the method comprising the steps of:

- 5 assigning a physical address to at least two of said ports;
- assigning a logical address to each said physical address;
- storing said logical address in a memory;
- defining each of said processing nodes to correspond to no more than one partition,
and;
- routing data within the system to correspond with the addresses of the resources
defined as a partition.

10 12. The method of claim 11 wherein one or more of said processing nodes may be
defined as corresponding to no partition.

13. The method of claim 11 further including the step of labeling data requests within
the system to a logical address corresponding to the partition issuing said request.

15 14. In a multiprocessor computer system comprising a plurality of processing nodes; a
shared, distributed system memory; and a communication pathway connecting said processing
nodes; wherein each one of said processing nodes includes at least one processor; and a
portion of said shared system memory coupled to said processor and said communication
pathway; said communications pathway comprised of communications ports each dedicated to
communicating with one of said processing nodes; a method for partitioning the resources of the
20 system into two or more partitions; the method comprising the steps of:

- assigning a physical address to each of said ports;
- assigning a logical address to each said physical address;
- storing said logical address in a memory;
- defining each logical address to correspond to no more than one partition;

assigning identification to each data transaction within the system which corresponds to the logical address to which the transaction is associated;

comparing said identified transaction with said logical address in memory corresponding to only one partition; and

5 routing data within the system to said ports defined as a member of the partition with which the data is associated.

15. The method of claim 14 further including the step of routing data requests within the system to a cache memory within said communications pathway with a logical address corresponding to the partition issuing said request, with said cache memory physically interconnected only with a port assigned to a physical address.

10